

Bidding for Army Career Specialties: Improving the ROTC Branching Mechanism*

Tayfun Sönmez[†]

September 2011

Abstract

Motivated by historically low retention rates of graduates at USMA and ROTC, the Army recently introduced *branch-for-service* incentives programs where cadets could bid an additional three years of active duty service obligation to obtain higher priority for their desired career specialties. The full potential of this highly innovative program is not utilized, due to the ROTC's choice of a poorly behaved cadet-branch matching mechanism. Not only does the ROTC mechanism effectively block the access of a large fraction of moderately high-skilled cadets to key career branches, but it is also highly vulnerable to preference manipulation and encourages effort reduction, potentially compromising human capital accumulation of the Army. Building on recent advances in matching markets, we propose a design that eliminates each of these deficiencies and also benefits the Army by mitigating several policy problems that the Army has identified. In contrast to the ROTC mechanism, our design utilizes market principles more elaborately, and it can be interpreted as a hybrid between a market mechanism and a priority-based allocation mechanism.

JEL Classification Numbers: C78, D63, D78

*I thank John Hatfield, Onur Kesten and Scott Kominers for helpful comments.

[†]Department of Economics, Boston College, Chestnut Hill, MA 02467. Phone: (617) 552-3690. Fax: (617) 552-2308. E-mail: sonmezt@bc.edu.

1 Introduction

In the last decade there has been a lot of activity and excitement among economists working on matching markets, a field that dates back to the seminal contribution of Gale and Shapley (1962). Theory matured to a point where matching theorists could make policy suggestions in key areas including education and health care. Research on the assignment of students to schools introduced by Balinski and Sönmez (1999) and further explored in the context of school choice by Abdulkadiroğlu and Sönmez (2003) resulted in the reform of student assignment mechanisms in Boston and New York City.¹ Research introducing kidney exchange to market design by Roth, Sönmez, and Ünver (2004, 2005, 2007) resulted in kidney exchange programs throughout the world including in the U.S., U.K., and the Netherlands. In his recent Congress testimony, Dr. Myron Gutmann (Assistant Director, Social, Behavioral, and Economic Sciences, NSF) emphasizes that research on matching markets has resulted in measurable gains for the U.S. taxpayer.²

The field of matching markets owes its recent success to the discovery of important practical applications backed by solid theory. In this paper, along with a companion paper (Sönmez and Switzer 2011), we introduce a brand new practical application of matching markets: *Cadet-branch matching* for U.S. Army programs. The two main programs the U.S. Army relies on to recruit officers are the United States Military Academy (USMA) and Reserve Officer Training Corps (ROTC). Graduates of USMA and ROTC enter active duty for an initial period of obligatory service upon completing their programs. The Active Duty Service Obligation (ADSO) is five years for USMA graduates, four years for ROTC scholarship graduates, and three years for ROTC non-scholarship graduates. Upon completion of this obligation, an officer may apply for voluntary separation or continue on active duty. Given the significant investment made for these officers, the Army has a strong preference for them to serve beyond their obligatory service. The low retention rate of these company-grade officers has been a major

¹See Abdulkadiroğlu, Pathak, Roth, and Sönmez (2005) and Abdulkadiroğlu, Pathak, and Roth (2005) for specifics in Boston and New York City respectively.

²Testimony available at www.nsf.gov/about/congress/112/mg_sberesearch_110602.jsp. Last accessed 09/08/2011.

issue for the U.S. Army since the late 1980s. In the last few years, the Army has responded to this challenge with unprecedented retention incentives, including *branch-for-service* incentives programs offered by both USMA and ROTC (Wardynski, Lyle, and Colarusso 2010).

During the fall semester of their senior year, USMA and ROTC cadets “compete” for their branch choices. The outcome of this branching process is essential for cadets not only because it determines their future specialties in the Army, but also because career advancement possibilities vary widely across different branches. There has been a long tradition of assigning branches to cadets based on their preferences and their merit ranking. This merit ranking is known as the *order-of-merit list* (OML) in the military and is based on a weighted average of academic performance, physical fitness test scores, and military performance. Until 2006, cadet-branching was an application of possibly one of the most straightforward resource allocation problems, and it was solved with a simple mechanism: The top OML cadet was assigned his first choice, the next cadet was assigned his top choice among remaining slots, and so on. This natural mechanism is known as *simple serial dictatorship* and was the mechanism of choice at both USMA and ROTC until 2006. Both programs changed their cadet-branching mechanisms that year in response to historically low retention rates of their graduates.

The idea behind this change was simple: Since branch choice is essential for most cadets, why not allow them to bid an additional period of ADSO for their desired branches? As part of the Army’s broader incentives program to combat the high attrition rate, cadets could “buy” priorities at a fraction of slots by agreeing to serve an additional three years of active service. The fraction of slots up for bidding is 25 percent for USMA, and 50 percent for ROTC. The new matching process is referred as the *branch-for-service* program for both USMA and ROTC, although the specifics of the new mechanisms are very different at the two programs. Sönmez and Switzer (2011) show that while the new mechanism has several shortcomings for USMA, a relatively easy fix that preserves all main aspects of the USMA design is available once cadet-branch matching is related to a recent important model introduced by Hatfield and Milgrom (2005) and further developed by Hatfield and Kojima (2008, 2010). In this paper, we show that the situation is rather different for ROTC, and a more substantial intervention is needed to design a fully satisfactory mechanism. We propose a design that relies more heavily on market

principles and show that not only are the shortcomings of the ROTC mechanism eliminated, but this new mechanism also mitigates several policy problems the Army has identified.

At first sight the shortcomings of the ROTC mechanism and the USMA mechanism appear to be very similar. They can both yield unfair outcomes where higher OML priority cadets can be envious of lower OML priority cadets even when they are willing to pay the increased cost; both mechanisms are vulnerable to preference manipulation, making the branch-selection a high stakes game for cadets; and both mechanisms can penalize cadets who improve their OML standing. What makes the ROTC mechanism a challenge to fix is a skill-based *affirmative action constraint* and the direct method ROTC has chosen to address it. Leadership at ROTC wants to avoid a situation where cadets of high skill are all concentrated in a few popular branches. To reach that objective they fully block the access of all cadets in the upper-half of OML to the last 35 percent of slots at each branch. Since at each branch only the last 50 percent of slots are available for the additional bid, this method ROTC uses means that cadets at the upper-half of the OML are, to a large extent, excluded from the branch-for-service program. While this had no impact for roughly the top 20 percent of cadets in the last three years, its adverse impact has been significant for cadets between the 20th and 50th percentiles. For each of the seven to eight popular branches, this policy resulted in what is referred as *dead zones* in Army jargon, to the severe detriment of a large fraction of cadets. In 2011, the dead zone affected cadets between the 20th and 50th percentiles for the most popular branch, aviation. No cadet in this range had access to aviation slots whether they were willing to pay the additional cost or not. In contrast, cadets from the 50th to 70th percentiles had access to all branches, including aviation, provided that they were willing to pay the extra cost. And not surprisingly, the closer the cadet was to the 50th percentile mark, the more compromised he was. For example, while cadets between the 20th and 30th percentiles faced only one dead zone in 2011, cadets between the 40th and 50th percentiles faced six to eight dead zones, depending on their their rank (see Figure 1, borrowed from an April 11, 2011, ROTC Accessions Process presentation³).

³Available at <http://www.purdue.edu/armyrotc/currentcadets/currentcadets.php> under *Accessions Briefing* heading. Last accessed at 8/31/2011.

Remarkably, it is possible to maintain the above-mentioned discontinuous ROTC priority structure and still fix the vulnerability of the ROTC mechanism to preference manipulation. Strategic simplification of the branching process is a major benefit to cadets, but we will later elaborate why it is also an essential improvement for the Army. That being said, we feel that to fix only the lack of incentive compatibility would be to do a partial job. The current ROTC priority structure is not compatible with the design of a fully satisfactory mechanism since it relies on the creation of the above-mentioned dead zones. So the key question is whether it is possible to implement the Army’s distributional goals without creating dead zones. We argue that the answer to this key question is affirmative although it will require that the Army rely more heavily on market principles. Here is our proposal: Currently cadets can bid a one-time bid of 3 years for the last 50 percent of the slots at each branch. If the Army allows cadets to bid more than 3 years, the role of the OML decreases and the role of willingness to serve increases in branch assignment. The idea is that once the highest possible bid is sufficiently high, motivated cadets in the lower-half of the OML will be able to outbid their less-motivated peers in the upper-half of the OML. Similarly, increasing the fraction of slots up for bidding would decrease the role of the OML and increase the role of willingness to serve in branch assignment. The target of assigning at least 35 percent of slots in each branch to cadets in the lower-half of the OML can be achieved by a mixture of these two adjustments. A reasonable starting point might be allowing 70 percent of the slots up for bidding at each branch and setting the maximum bid at a relatively high price, say 9 years of extra active duty service obligation. We refer to the resulting priority structure as *Bid-for-Your-Career* (BfYC) priorities.

One might argue that this more elaborate reliance on market principles further undermines the role of the order-of-merit system in cadet branching. We believe that is not the case. The creation of dead zones by ROTC priorities already severely undermines the role of the order-of-merit system. Under our proposed BfYC priorities, the sharp discontinuity created by ROTC priorities will be avoided, and rather than favoring arbitrary cadets in the lower-half of the OML, those who are most willing to serve will be favored.

Finding an “indirect” way to implement the Army’s distributional goals is the most challenging part of the design. Once this hurdle is cleared, recent advances in theoretical matching

literature by Hatfield and Milgrom (2005), and Hatfield and Kojima (2008, 2010) give us a lot of mileage to design a mechanism that eliminates the above-mentioned shortcomings of the ROTC mechanism: A *cadet-optimal stable mechanism* (COSM) is well-defined under BfYC priorities, it is stable, and it Pareto dominates any other stable mechanism. This mechanism always yields a fair outcome in the sense that a high-priority cadet never envies the full assignment of a lower-priority cadet (although he could be envious of the branch portion of the assignment). Truth-telling is always optimal under this mechanism (i.e. COSM is strategy-proof), and an increase in OML standing never hurts a cadet which brings cadets' incentives in line with the hard work necessary for their academic, physical fitness, and military studies. For the basic case where no slot is reserved for bidding, this mechanism reduces to the *Gale-Shapley agent-optimal stable mechanism*, recently adopted by Boston Public Schools (BPS) for assignment of K-12 students to public schools and by the New York City Department of Education for assignment of high school students to public high schools.⁴ The desire to replace highly manipulable student assignment mechanisms with their strategy-proof counterparts was one of the key reasons for these reforms in both school districts, and especially at BPS.

The current ROTC mechanism is highly deficient from a mechanism design perspective, and replacing it with COSM under BfYC priorities has numerous benefits for cadets, as we have presented. This potential reform will also benefit the Army on a number of important policy issues. As previously mentioned, the Army has a major attrition problem and branch-for-service incentives were adopted as a response to this problem. Using a mechanism that explicitly singles out cadets from the 20th to 50th percentiles might frustrate cadets in this reasonably high-skill cadet group, thereby increasing its attrition rate. In contrast, COSM under BfYC priorities favors cadets who are most willing to serve, thus increasing the cost of leaving the Army right after the base active duty service obligation. By allowing cadets to bid more than three years, the mechanism is also likely to significantly boost the man-year gains of the mechanism. Another major benefit of adopting COSM (even if ROTC priorities are maintained) has to do with its strategy-proofness. A recent study by Lim et al. (2009) investigates the cause of highly

⁴For the case of a uniform priority list across all schools/branches, the mechanism further reduces to simple serial dictatorship.

undesired minority underrepresentation in leadership ranks of the Army. They observe that while 80 percent of generals are from *combat arms branches*, minorities do not target these key career branches as much as their white peers. They also observe that minorities tend to have lower OML than their white peers, and thus the ROTC mechanism gives them an incentive to target less competitive branches. Hence they conclude that part of what seems to be a lack of interest on the part of minorities for combat arms branches might be strategic. They are unable to make policy suggestions since adequate remedies would depend on to what extent the lack of minority representation in key career fields is an artifact of preference manipulation. In simple terms, the Army cannot interpret preference data because its mechanism is highly manipulable. This is entirely avoided under COSM and will allow the Army to implement adequate policies to improve diversity in its senior ranks. Another benefit of our proposed reform to the Army is the flexibility of the leadership of each branch to determine its own base priority ranking. This flexibility, which is absent from the current ROTC mechanism, is highly desired for some branches such as military intelligence (Besuden II 2008). Finally, the current ROTC mechanism highly encourages cadets from the 20-50 percentile range of OML to reduce their ranking to avoid falling in dead zones in cadet-branching. The Army clearly benefits by adopting a mechanism that promotes high effort levels over slacking.

In addition to introducing a new practical market design problem, our paper, along with Sönmez and Switzer (2011), brings a new perspective to a recent debate in matching markets. The cadet-branch matching problem is a special case of the *matching with contracts* model (Hatfield and Milgrom 2005) although a *substitutes* condition that has been key to Hatfield and Milgrom's analysis is not satisfied in our context. The matching with contracts model owes much of its early success to the perception that it subsumes and unifies the Gale and Shapley (1962) college admissions model and the Kelso and Crawford (1982) labor market model, among others.⁵ In a highly surprising result, Echenique (2011) has shown that under the substitutes condition, the matching with contracts model can be embedded within the Kelso and Crawford (1982) labor market model, thus showing that the two models are isomorphic.⁶ As emphasized

⁵Kelso and Crawford (1982) build on the analysis of Crawford and Knoer (1981).

⁶See also Kominers (2011) for an extension of this isomorphism to many-to-many matching.

by Echenique (2011), the substitutes condition is key for this isomorphism. In particular, that paper indicates that a recent theory paper by Hatfield and Kojima (2010) analyzes matching with contracts under weaker conditions, and his embedding does not work under their conditions. Although Hatfield and Kojima (2010) do not offer any applications under their weaker *unilateral substitutes* condition, they show that a number of key results on the *agent-optimal stable mechanism* persist under this condition. Remarkably, although the substitutes condition fails in our framework, the unilateral substitutes condition is satisfied. Thus cadet-branch matching is the first practical application of the full generality of the matching with contracts model.

2 The Model

Since 2006, both USMA and ROTC cadets are given an option to sign one or more *branch-of-choice* contracts that increase their priorities at branches of their choosing in exchange for three additional years of active military service. A branch-of-choice contract does not guarantee that a cadet will receive a slot at a branch, nor does signing one necessarily oblige him for the three additional years of service even if the cadet is assigned a slot at the branch. Loosely speaking, ROTC has changed the priorities for the last 50 percent of the slots at each branch and has given priority to cadets who have committed to three additional years of active service.⁷ A cadet who signs a branch-of-choice contract is obliged to serve the additional three years of service only if he receives a slot from the last 50 percent of the capacity of a branch.

Prior to adoption of branch-for-service programs, cadet-branch matching was an application of the most basic form of what is known as the *student placement problem* in the literature (Balinski and Sönmez 1999). Since the adoption of branch-for-service programs, the nature of the cadet-branch matching problem has changed in two important ways. First of all, the outcome of the problem is no longer merely an assignment of branches to cadets but rather an assignment of branches along with the terms of these assignments. And second, willingness to pay a “higher price” started playing a role in determining who has higher claims to a fraction

⁷For the case of USMA, priorities are changed only at the last 25 percent of the slots.

of the slots.

We are now ready to formally introduce the problem.

A **cadet-branch matching problem** consists of

1. a finite set of cadets $I = \{i_1, i_2, \dots, i_n\}$,
2. a finite set of branches $B = \{b_1, b_2, \dots, b_m\}$,
3. a vector of branch capacities $q = (q_b)_{b \in B}$,
4. a set of “terms” $T = \{t_1, \dots, t_k\}$,
5. a list of cadet preferences $P = (P_i)_{i \in I}$ over $(B \times T) \cup \{\emptyset\}$, and
6. a list of base priority rankings $\pi = (\pi_b)_{b \in B}$.

Here $t \in \mathbb{R}_+$ for each $t \in T$ with $t_1 < t_2 < \dots < t_k$, and a cadet who is assigned the pair (b, t) commits to serving in the military for at least t years. For any branch b , the function $\pi_b : I \rightarrow \{1, \dots, n\}$ represents the base priority ranking of cadets for branch b , and $\pi_b(i) < \pi_b(j)$ means that cadet i has higher claims to a slot at branch b than cadet j , other things being equal. Throughout the paper we fix the set of cadets I , the set of branches B , the vector of capacities q , and the set of terms T . Hence each problem is defined by a preference list along with a base priority list.

We assume that cadet preferences are strict and are such that for any cadet i , any branch b , and any pair of terms t, t' ,

$$t < t' \Leftrightarrow (b, t) P_i (b, t').$$

Hence cadet preferences are assumed to be **monotonically decreasing** in length of service. We also assume that preferences over branches are independent of the service obligation, and thus each cadet has well-defined preferences over branches.⁸ Let \succ_i denote cadet preferences

⁸While both these assumptions are rather natural in the present context, a careful reader will observe that neither is essential to our analysis. Both assumptions are made for the ease of exposition. Indeed, assuming monotonicity makes the proof of our uniqueness result in Proposition 4 marginally more challenging since it forces the constructed preferences to satisfy this natural requirement. The description and evaluation of the

over branches alone. For any cadet i , any pair of branches b, b' , and any term t , we have

$$b \succ_i b' \Leftrightarrow (b, t) P_i(b', t).$$

Let \mathcal{P} denote the set of all preferences over $(B \times T) \cup \{\emptyset\}$, and \mathcal{Q} denote the set of all preferences over B .

A **contract** $x = (i, b, t) \in I \times B \times T$ specifies a cadet i , a branch b , and the terms of their match. Let $X \equiv I \times B \times T$ be the set of all contracts. Given a contract $x = (i, b, t)$, let $x_I = i$ denote the cadet, $x_B = b$ denote the branch, and $x_T = t$ denote the terms of the contract x .

An **allocation** $X' \subset X$ is a set of contracts such that each cadet appears in at most one contract and no branch appears in more contracts than its capacity. Let \mathcal{X} denote the set of all allocations.

Given a cadet $i \in I$ and an allocation $X' \subset X$ with $(i, b, t) \in X'$, let $X'(i) = (b, t)$ denote the **assignment** of cadet i under allocation X' . If a cadet i remains unmatched under allocation X' , then $X'(i) = \emptyset$.

For a given problem, an allocation X' is **fair** if $X'(j) P_i X'(i) \Rightarrow \pi(j) < \pi(i)$ for any pair of cadets i, j . That is, a higher-priority cadet can never envy the assignment of a lower-priority cadet under a fair allocation. Note that it is still possible for a higher-priority cadet to envy the branch assigned to a lower-priority cadet under a fair allocation. Consider a high-priority cadet i with an assignment $X'(i) = (b, t)$ and a low-priority cadet j with $X'(j) = (b', t')$ with $t' > t$. While fairness rules out $X'(j) P_i X'(i)$, it is still possible that $b(j) \succ_i b(i)$. A low-priority cadet may be able to get a more preferred branch, because he is willing to pay a higher price for it.

A **mechanism** is a **strategy space** S_i for each cadet i along with an **outcome function** $\varphi : (S_1 \times S_2 \times \cdots \times S_n) \rightarrow \mathcal{X}$ that selects an allocation for each strategy vector $(s_1, s_2, \dots, s_n) \in (S_1 \times S_2 \times \cdots \times S_n)$. Given a cadet i and strategy profile $s \in S$, let s_{-i} denote the strategy of all cadets except cadet i . A mechanism is **fair** if it always selects a fair allocation.

ROTC mechanism also becomes somewhat more convenient when we assume that preferences over branches are independent of service obligation, for otherwise a cadet might not have well-defined preferences over branches alone.

A **direct mechanism** is a mechanism where the strategy space is simply the set of preferences \mathcal{P} for each cadet i . Hence a direct mechanism is simply a function $\varphi : \mathcal{P}^n \rightarrow \mathcal{X}$ that selects an allocation for each preference profile.

A highly desirable property of a direct mechanism is that it is always in cadets' best interests to be entirely truthful about their preferences. Hence, cadets can never benefit from “gaming” such mechanisms.

Definition: A direct mechanism φ is **strategy-proof** if

$$\varphi(P_i, P_{-i}) \succeq_i \varphi(P'_i, P_{-i}) \text{ for any } i \in I, P_{-i} \in \mathcal{P}^{n-1} \text{ and } P_i, P'_i \in \mathcal{P}.$$

That is, no matter which cadet i we consider, no matter what his true preferences P_i are, no matter which preferences P_{-i} the rest of the cadets report (true or not), and no matter which potential “misrepresentation” P'_i cadet i considers, truthful preference revelation is in his best interests.

One of the most important parameters of the cadet-branch matching problem is base branch priorities π . Clearly, a reasonable mechanism would not penalize a cadet because of an improvement of his base priorities. Given a pair of base priority rankings π_b^1, π_b^2 , we will say that π_b^1 is an **unambiguous improvement** for cadet i over π_b^2 at branch b if

1. the relative ranking between all cadets except cadet i remains exactly the same between π_b^1 and π_b^2 , although
2. the standing of cadet i is strictly better under π_b^1 than under π_b^2 .

Given two lists of base priority rankings π^1, π^2 , we will say that π^1 is an **unambiguous improvement** for cadet i over π^2 if

1. either $\pi_b^1 = \pi_b^2$ or π_b^1 is an **unambiguous improvement** for cadet i for any branch b , and
2. π_b^1 is an **unambiguous improvement** for cadet i over π_b^2 for some branch b .

Definition: A mechanism **respects improvements** if a cadet never receives a strictly worse assignment as a result of an unambiguous improvement in his priority ranking.

Observe that the failure of this property hurts the mechanism not only from a normative perspective, but also via the adverse incentives it creates in case cadet effort plays any role in determining the base priorities. As in most merit-based resource allocation problems, this is the case for both USMA and ROTC cadet branching.

3 ROTC Cadet-Branch Matching

Both USMA and ROTC have branch-for-service incentives programs where cadets can sign branch-of-choice contracts for one or more branches extending their active duty service obligation for three years in exchange for increased priority at these branches. Sönmez and Switzer (2011) analyze the USMA mechanism and show that while it has several deficiencies, these can be overcome by a mechanism that is remarkably similar to the USMA mechanism. The situation is somewhat more involved for the ROTC mechanism, and an affirmative action constraint for lower-performance cadets makes the design of a fully satisfactory mechanism a challenge. Since 2006, $T = \{t_1, t_2\}$ for the case of ROTC branching, and in this section we refer to t_1 as the **base cost**, t_2 as the **increased cost**, and any contract with increased costs t_2 as a **branch-of-choice contract**. We are now ready to describe the ROTC mechanism.

ROTC Mechanism

All cadets are ranked by a single *order-of-merit* list ranking π^{OML} that is based on a weighted average of academic performance, physical fitness test scores, and military performance. Since the adoption of the ROTC branch-for-service incentives program, $\pi_b = \pi^{OML}$ for any branch b under ROTC base priorities. Under the ROTC mechanism, each cadet submits a preference ranking of branches \succ'_i and he can sign a branch-of-choice contract for any of his top three choices under \succ'_i . Let B_i denote the (possibly empty) set of branches for which cadet i signs a branch-of-choice contract.

For a given order-of-merit list π^{OML} and a strategy-profile $(\succ'_i, B_i)_{i \in I}$, the outcome of the ROTC mechanism is obtained as follows:

Consider each cadet one - at - a - time, following their priority order. The treatment of cadets at the top 50 percent of the OML is slightly different from that of those at the bottom

50 percent.

- For each cadet at the top 50 percent of the OML, consider the following six options in the given order, and if none of them works, leave the cadet unassigned.
 1. Assign the cadet his first-choice branch at base cost t_1 , if less than 50 percent of the slots are full at his first choice.
 2. Assign the cadet his first-choice branch at increased cost t_2 , if he signed a branch-of-choice contract and less than 65 percent of the slots are full at his first choice.
 3. Assign the cadet his second-choice branch at base cost t_1 , if less than 50 percent of the slots are full at his second choice.
 4. Assign the cadet his second-choice branch at increased cost t_2 , if he signed a branch-of-choice contract and less than 65 percent of the slots are full at his second choice.
 5. Assign the cadet his third-choice branch at base cost t_1 , if less than 50 percent of the slots are full at his third choice.
 6. Assign the cadet his third-choice branch at increased cost t_2 , if he signed a branch-of-choice contract and less than 65 percent of the slots are full at his first choice.
- For each cadet at the bottom 50 percent of the OML, consider the following six options in the given order, and if none of them works, leave the cadet unassigned.
 1. Assign the cadet his first-choice branch at base cost t_1 , if less than 50 percent of the slots are full at his first choice.
 2. Assign the cadet his first-choice branch at increased cost t_2 , if he signed a branch-of-choice contract and not all slots are full at his first choice.
 3. Assign the cadet his second-choice branch at base cost t_1 , if less than 50 percent of the slots are full at his second choice.
 4. Assign the cadet his second-choice branch increased higher cost t_2 , if he signed a branch-of-choice contract and not all slots are full at his second choice.

5. Assign the cadet his third-choice branch at base cost t_1 , if less than 50 percent of the slots are full at his third choice.
6. Assign the cadet his third-choice branch at increased cost t_2 , if he signed a branch-of-choice contract and not all slots are full at his first choice.

If a cadet remains unassigned under the ROTC mechanism, his branch assignment is determined by the Department of the Army Branching Board.

The structure of branch priorities at ROTC is in conflict with the design of a fully satisfactory mechanism. For each branch b , **ROTC branch priorities** are given as follows:⁹

- For the top 50 percent of the slots, the priority is based on cadet OML.¹⁰
- The next 15 percent of the slots are *reserved* for cadets who have signed a branch-of-choice contract for branch b , and among them priority is based on cadet OML.
- The last 35 percent of the slots are *reserved* for cadets who are at the bottom 50 percent of the OML who have signed a branch-of-choice contract for branch b . Among them priority is based on cadet OML.

Note that there is an “affirmative action” constraint for the last 35 percent of the slots at each branch, and cadets at the upper half of the OML ranking are denied access to these slots whether they are willing to pay the increased cost or not. Given these branch priorities, the ROTC mechanism clearly fails to be fair, nor does it respect improvements. For a given branch, the range of the OML where higher-ranking cadets lose priority to cadets in the lower-half of the OML is referred as the **dead zone** by the Army. In 2011, eight of the most popular branches had a dead zone (see Figure 1). These branches and their dead zones are:

1. Aviation with cadets between 20-50 percent of the OML,
2. Infantry with cadets between 30-50 percent of the OML,

⁹Since 2010, priorities for Engineering is slightly different from other branches: Half of the slots for Engineering are reserved for cadets with Engineering degrees, but otherwise the same priority structure is followed.

¹⁰The first choice is guaranteed for each cadet at the top 10 percent of the OML.

3. Medical Service with cadets between 31-50 percent of the OML,
4. Armor with cadets between 35-50 percent of the OML,
5. Engineering with cadets between 38-50 percent of the OML,
6. Military intelligence with cadets between 40-50 percent of the OML,
7. Military police with cadets between 43-50 percent of the OML, and
8. Finance with cadets between 47-50 percent of the OML.

In contrast to unfortunate cadets “falling in” one or more dead zones, cadets in the range 50-70 percent of OML had full access to each of the 16 branches provided they had signed a branch-of-choice contract.

It is also easy to see that most cadets need to strategize well under the ROTC mechanism. The most obvious reason for this is that the mechanism only considers the top three branch choices, and given the high stakes, cadets need to choose these three branches wisely. The choice of branch-of-choice contracts is not an easy task either, since a branch at increased cost is considered right after the same branch at base cost. In Section 6.2 we discuss in detail why the vulnerability of the ROTC mechanism to preference manipulation is a major obstacle for the U.S. Army in the analysis of a highly debated policy issue.

Before we propose an alternative mechanism that corrects each of these shortcomings, we need to relate ROTC cadet-branch matching to a recent important model.

4 Matching with Contracts

The cadet-branch matching problem is a special case of the **matching with contracts** model (Hatfield and Milgrom 2005). In the original Hatfield-Milgrom model, each branch (hospitals in their framework) has preferences over sets of agent-cost pairs. These hospital preferences induce a *choice set* from each set of contracts, and it is this choice set (rather than hospital preferences) that is key in the model. In the present framework, branches are not agents and

they do not have preferences. However, branches have priorities over cadet-cost pairs, and these priorities also induce choice sets. This is the sense in which the cadet-branch matching problem is a special case of matching with contracts.

We next present a few key concepts from matching with contracts before we propose a new mechanism for ROTC cadet branching. Recall that each cadet i has preferences P_i over all branch-cost pairs. Equivalently, each cadet i has strict preferences over all contracts that include him. A cadet may not be assigned a branch-cost pair under the ROTC mechanism. In that case his assignment is determined by the Department of the Army Branching Board (DABB) and he is only charged the base cost t_1 . Hence, in our model, a cadet who is assigned \emptyset by the ROTC mechanism (or any alternative mechanism) receives a branch that is determined by DABB at the base cost t_1 . This “lottery” may be more preferred for some cadets than a number of assignments (such as assignments at high cost or assignments with highly undesirable branches). A contract (i, b, t) is **undesired** for cadet i if $\emptyset P_i(b, t)$.

Given a set of contracts X' and a preference relation P_i , define the **choice of cadet i** from X' , $C_i(X')$, to be \emptyset if all contracts in X' that include cadet i are undesired and to be the singleton that consists of the most preferred contract of cadet i in X' under P_i otherwise. That is,

$$C_i(X') \equiv \begin{cases} \emptyset & \text{if } \nexists (i, b, t) \in X' \text{ such that } (b, t) P_i \emptyset, \\ (i, b, t) & \text{if } (b, t) P_i \emptyset \text{ and } (b, t) P_i (b', t') \text{ for any } (i, b', t') \in X' \setminus \{(i, b, t)\}. \end{cases}$$

In general, the **choice of branch b** from a set of contracts X' , $C_b(X')$, depends on the policy on who has higher claims for slots at branch b . We refer $R_b(X') \equiv X' \setminus C_b(X')$ as the **rejected set**.

We next describe the **choice of branch b under ROTC priorities**: Let $T = \{t_1, t_2\}$ where t_1 is the base cost and t_2 is the increased cost. Given a set of contracts X' , choice of branch b under ROTC priorities is obtained as follows:

Phase 0: Remove all contracts that involve another branch b' and add them all to the rejected set $R_b(X')$. Hence each contract that survives Phase 0 involves branch b .

Phase 1.1: For the first $0.5q_b$ potential elements of $C_b(X')$, simply choose the contracts

with highest OML-priority cadets one at a time. When two contracts of the same cadet are available, choose the contract with the base cost t_1 and reject the other one, including it in $R_b(X')$. Continue until either all contracts are considered or $0.5q_b$ elements are chosen for $C_b(X')$. If the former happens, terminate the procedure, and if the latter happens, proceed with Phase 1.2.

Phase 1.2: Remove all surviving contracts with base cost t_1 and add them all to the rejected set $R_b(X')$. Proceed with Phase 2.1 if there is at least one surviving contract and terminate the procedure otherwise.

Phase 2.1: All remaining contracts have increased cost t_2 . Among them include in $C_b(X')$ the contracts with highest OML-priority cadets for the next $0.15q_b$ potential elements of $C_b(X')$. Continue until either all contracts are considered in X' or $0.65q_b$ elements are chosen for $C_b(X')$. For the former case terminate the procedure. For the latter case, terminate the procedure if all contracts in X' are considered, and proceed with Phase 2.2 otherwise.

Phase 2.2: Remove all surviving contracts that belong to cadets from the upper half of the OML list π^{OML} and add them all to the rejected set $R_b(X')$. Proceed with Phase 3 if there is at least one surviving contract and terminate the procedure otherwise.

Phase 3: All remaining contracts have increased cost t_2 and belong to cadets from the lower half of the OML list π^{OML} . Among them include in $C_b(X')$ the contracts with highest OML-priority cadets for the last $0.35q_b$ potential elements of $C_b(X')$. Reject all remaining contracts and terminate the procedure.

Definition: Priorities are **fair** if for any branch b the induced choice function C_b is such that, for any set of contracts X' and any pair of contracts $x, y \in X'$ with $x_B = y_B = b$,

$$\left. \begin{array}{l} \pi_b(y_I) < \pi_b(x_I), \\ y_T = x_T, \text{ and} \\ x \in C_b(X') \end{array} \right\} \implies \exists z \in C_b(X') \text{ such that } z_I = y_I.$$

That is, if a contract x of a lower-priority cadet is chosen, then a contract z of a higher-priority cadet who is willing to pay as much under a reference contract y shall also be chosen under fair

priorities. Here the chosen contract of cadet y_I can be the reference contract y or an alternative contract z .

Observation: Because of Phase 2.2, ROTC priorities are not fair. Cadets from the upper half of the OML are simply denied for the last 35 percent of slots at each branch.

Since the seminal paper of Gale and Shapley (1962), a condition known as stability has been central to the analysis of two-sided matching markets as well as the allocation of indivisible goods based on priorities.¹¹ The matching with contracts model has also evolved around the stability axiom. Formally, an allocation X' is **stable** if

1. $\bigcup_{i \in I} C_i(X') = X'$,
2. $\bigcup_{b \in B} C_b(X') = X'$, and
3. there exists no cadet i , branch b , and contract $x = (i, b, t) \in X \setminus X'$ such that

$$\underbrace{\{x\} = C_i(X' \cup \{x\})}_{\text{(i.e. } (b,t)P_i X'(i)\text{)}} \text{ and } x \in C_b(X' \cup \{x\}).$$

In the context of cadet-branch matching, the only plausible allocations are the stable ones. Note that if the first requirement fails then there is a cadet who prefers to reject a contract that involves him (or equivalently, there is a cadet upon whom is imposed an undesired contract); if the second requirement fails then there exists a branch that would rather reject one of its contracts; and if the third requirement fails then there exists an unselected contract (i, b, t) where not only cadet i prefers pair (b, t) to his assignment, but also contract x has sufficiently high priority to be selected by branch b .

The following two properties of branch priorities have played an important role in the analysis of matching with contracts:

Definition: Priorities satisfy the **law of aggregate demand** for branch b if $X' \subset X'' \Rightarrow |C_b(X')| \leq |C_b(X'')|$.

¹¹See Roth and Sotomayor (1990) and Sönmez and Ünver (2010) for comprehensive surveys on the role of stability in two-sided matching markets and the allocation of indivisible goods based on priorities.

That is, the size of the choice set never shrinks as the set of contracts grows under the law of aggregate demand.

Definition: Elements of X are **substitutes** for branch b under a choice function C_b if for all $X' \subset X'' \subseteq X$, we have $R_b(X') \subseteq R_b(X'')$.

That is, contracts are substitutes if a contract that is chosen from a larger set X'' is chosen from any of its subsets $X' \subset X''$ as well. Equivalently, any contract that is rejected from a smaller set X' is also rejected from any larger set X'' that contains X' . If elements of X are substitutes, then the set of stable allocations is non-empty (Hatfield and Milgrom 2005). Indeed, until Hatfield and Kojima (2008) showed otherwise, it was thought to be a necessary condition for the guaranteed existence of a stable allocation.

It is easy to see that ROTC priorities satisfy the law of aggregate demand but not the substitutes condition: Consider a cadet i who is at the lower-half of the OML. Between contracts $x = (i, b, t_1)$ and $y = (i, b, t_2)$, the cheap contract x might be chosen while the expensive one y is rejected from a small set of contracts X' with little competition, although the choice is reversed for $X'' \supset X'$ where competition for slots is higher.

A weaker condition, recently introduced by Hatfield and Kojima (2010), is as follows:

Definition: Elements of X are **unilateral substitutes** for branch b if, whenever a contract $x = (i, b, t)$ is rejected from a smaller set X' *even though x is the only contract in X' that includes cadet i* , contract x is also rejected from a larger set X'' that includes X' .

We will say that **priorities satisfy the unilateral substitutes condition** when elements of X are unilateral substitutes under these priorities.

Lemma 1 *ROTC priorities satisfy the unilateral substitutes condition and the law of aggregate demand.*

4.1 Cumulative Offer Algorithm

We are ready to introduce the **cadet-optimal stable mechanism (COSM)**, which is simply a natural extension of the celebrated **agent-optimal stable mechanism** (Gale and Shapley 1962), to matching with contracts:

The strategy space of each cadet is \mathcal{P} under the COSM, and hence it is a direct mechanism. Fix a choice function C_b for each branch b . Given a preference profile $P \in \mathcal{P}$, the following algorithm can be used to find the outcome of COSM.

Cumulative Offer Algorithm¹²

Step 1: Start the offer process with the highest OML priority cadet $\pi^{OML}(1) = i(1)$. Cadet $i(1)$ offers his first-choice contract $x_1 = (i(1), b(1), t)$ to branch $b(1)$ that is involved in this contract. Branch $b(1)$ holds the contract if $x_1 \in C_{b(1)}(\{x_1\})$ and rejects it otherwise. Let $A_{b(1)}(1) = \{x_1\}$ and $A_b(1) = \emptyset$ for all $b \in B \setminus \{b(1)\}$.

Step 2: Let $i(2)$ be the highest OML priority cadet for whom no contract is currently held by any branch. Cadet $i(2)$ offers his most-preferred contract $x_2 = (i(2), b(2), t)$ that has not been rejected in the previous step to branch $b(2)$. Branch $b(2)$ holds the contract if $x_2 \in C_{b(2)}(A_{b(2)}(1) \cup \{x_2\})$ and rejects it otherwise. Let $A_{b(2)}(2) = A_{b(2)}(1) \cup \{x_2\}$ and $A_b(2) = A_b(1)$ for all $b \in B \setminus \{b(2)\}$.

In general, at

Step k : Let $i(k)$ be the highest OML priority cadet for whom no contract is currently held by any branch. Cadet $i(k)$ offers his most-preferred contract $x_k = (i(k), b(k), t)$ that has not been rejected in previous steps to branch $b(k)$. Branch $b(k)$ holds the contract if $x_k \in C_{b(k)}(A_{b(k)}(k-1) \cup \{x_k\})$ and rejects it otherwise. Let $A_{b(k)}(k) = A_{b(k)}(k-1) \cup \{x_k\}$ and $A_b(k) = A_b(k-1)$ for all $b \in B \setminus \{b(k-1)\}$.

The algorithm terminates when each cadet either has an offer that is on hold by a branch or has consumed all acceptable contracts. Since there is a finite number of contracts, the algorithm terminates after a finite number K of steps. All contracts held at this final Step K are finalized and the final allocation is $\bigcup_{b \in B} C_b(A_K)$.

Remark 1 *While the choice of the cadet making the offer at any given step is uniquely defined by the above-described cumulative offer algorithm, the same outcome is obtained regardless of the choice of cadet as long as there is no contract held by any branch that involves the cadet*

¹²The following description is borrowed mostly from Hatfield and Kojima (2010).

making the offer. Indeed, Hatfield and Kojima (2010) describe the algorithm without explicitly specifying the order of agents making offers.

Fix a choice function C_b for each branch b . Given a preference profile $P \in \mathcal{P}$, let $\varphi(P)$ denote the outcome of the COSM.

We will rely on the following pair of results by Hatfield and Kojima (2010):

Theorem 1 (Hatfield and Kojima 2010) *Suppose the priorities satisfy the unilateral substitutes condition. Then the cumulative offer algorithm produces a stable allocation. Moreover, this allocation is weakly preferred by any cadet to any stable allocation.*

Theorem 2 (Hatfield and Kojima 2010) *Suppose the priorities satisfy the unilateral substitutes condition and the law of aggregate demand. Then the induced COSM is strategy-proof.*

We are ready to present our first main result.

Proposition 1 *Suppose that the priorities satisfy the unilateral substitutes condition and the law of aggregate demand. Then the COSM is fair if and only if the priorities are fair.*

Let φ^{ROTC} be the cadet-optimal stable mechanism induced by the ROTC priorities. Our next result states that φ^{ROTC} fixes only some of the deficiencies of the ROTC mechanism. As we have argued before, ROTC priorities are not compatible with a fully satisfactory mechanism.

Proposition 2 *The outcome of φ^{ROTC} is stable under ROTC priorities and is weakly preferred by any cadet to any stable allocation. Moreover φ^{ROTC} is strategy-proof. However, it is neither fair nor respects improvements.*

With the introduction of branch-for-service incentives, USMA and ROTC adopted two different mechanisms, relying on redefined cadet claims on branches via USMA priorities and ROTC priorities respectively. While both new mechanisms suffer from similar deficiencies, a fix for the USMA mechanism is straightforward upon relating cadet branching to matching with contracts. That is mostly because the USMA priorities not only satisfy the unilateral

substitutes condition and the law of aggregate demand (Sönmez and Switzer 2011), they are also fair. On the other hand, while ROTC priorities satisfy the unilateral substitutes condition and the law of aggregate demand, they are unfair. Hence it is necessary to seek an alternative priority structure in order to design a satisfactory mechanism for ROTC branching.

As we discussed earlier, the range of the OML just above the 50 percent line is referred as the “dead zone” by the Army. The formal definition stated in an April 10, 2011, dated Accessions Briefing is as follows:

Dead Zone: The area on the branch bar graph where it is impossible for a cadet to receive a certain branch.

Recall that ROTC priorities fail to be fair because cadets in the upper-half of OML are denied the last 35 percent of slots at each branch. There is only one reason for this unusual choice of ROTC priorities. The Army desires to allocate skill somewhat evenly across its branches. The implications are substantial for cadets who face one or more dead zones. In the next section we propose an “indirect” approach to address this challenging distributive objective. The resulting mechanism will not only restore all flawed aspects of the ROTC mechanism, but also benefit the Army in a number of directions.

5 Bidding for Priorities

As we have argued, current ROTC priorities are not compatible with a fully satisfactory mechanism. This observation leads to the following natural question: Could it be possible to reach the Army’s distributional goal without creating a dead zone? We argue that the answer is yes. Our approach is based on increasing the highest price that cadets can bid for their desired branches and adjusting branch priorities in the following way:

- The top λ percent of the slots will be allocated following the OML; whereas
- cadets who are willing to serve in the military longer will have higher priority for the last $(1 - \lambda)$ percent of the slots.

There are currently two “prices” for branch assignment under the ROTC mechanism: the base price and the increased price that is three years in addition to the base price. Hence $T = \{t_1, t_2\}$ where $t_2 = t_1 + 3$ years under the current mechanism. Under our proposed mechanism the set of terms is larger and cadets are able to bid more than three years of increased service. In particular, we need the highest price to be relatively large, perhaps nine years of increased service, so that only the most motivated cadets will be willing to pay the highest cost.¹³ Observe that this will decrease the role of the OML and increase the role of willingness to serve in branch priorities. Another factor that will shift the balance in favor of willingness to serve is increasing the fraction of slots up for bidding. The idea is that the Army’s distributional goal of allocating skill reasonably evenly across branches can be achieved if the role of willingness to serve is sufficiently increased and the role of the OML is sufficiently decreased in branch priorities. Implicit here is the assumption that there is no strong positive correlation between the OML and willingness to serve. This is a sensible assumption, since Wardynski, Lyle, and Colarusso (2010) report that the retention rate is lower among higher-OML cadets. A reasonable starting point could be $\lambda = 30$ percent so that cadets can bid for 70 percent of the slots at each branch. Clearly the lower the parameter λ is, the higher the access of lower OML cadets for highly-sought branches, provided that they are willing to pay the price. In contrast to the current mechanism, there will not be any arbitrary dead zones and motivation to serve the Army will play a more important role in cadet branching. Our proposal builds on the following perspective offered by Wardynski, Lyle, and Colarusso (2010), who played a central role in the design of branch-for-service incentives programs:

The branch and post incentives also raised concerns. Devoted supporters of the ROTC and West Point Order of Merit (OML) system for allocating branches and posts objected that low OML cadets could buy their branch or post of choice ahead of higher OML cadets. Since branch and post assignments represent a zero sum

¹³It is important to emphasize that increasing the highest possible bid beyond three years is a feasible design consideration. In addition to branch-for-service incentives, the Army has a *post-for-service* incentives program as well as a *graduate school-for-service* incentives program, and cadets are allowed to apply for up to two of these programs for a total of six years of additional active duty service obligation.

game, the ability of cadets with a lower OML ranking to displace those above them was viewed by some as unfair or as undermining the OML system. However, rather than undermining the legacy system or creating inequities, the branch and post incentives program makes willingness to serve a measure of merit in branching and posting, thus providing taxpayers a fair return on their officer accessions investment.

For a given λ and set of terms $T = \{t_1, \dots, t_k\}$, the choice of branch b from a set of contracts X' is obtained as follows under our proposed **Bid-for-Your-Career (BfYC) priorities**.

Phase 0: Remove all contracts that involve another branch b' and add them all to the rejected set $R_b(X')$. Hence each contract that survives Phase 0 involves branch b .

Phase 1: For the first λ percent potential elements of $C_b(X')$, simply choose the contracts with highest OML-priority cadets one at a time. When multiple contracts of the same cadet are available, choose the contract with the lowest cost and reject the other ones, including them in $R_b(X')$. Continue until either all contracts are considered or λ percent of the capacity is full. If the former happens, terminate the procedure, and if the latter happens, proceed with Phase 2.

Phase 2: For the last $(1 - \lambda)$ percent potential elements of $C_b(X')$, choose the contracts with highest costs while using the OML to break ties. When multiple contracts of the same cadet are available, choose the contract with the highest cost and reject the other ones, including them in $R_b(X')$. Continue until either all contracts are considered or the capacity is full. Reject any remaining contracts.

Remark 2 *In order to deviate minimally from the current ROTC priorities, we defined BfYC priorities based on a fixed base priority ranking, the OML, for each branch. Clearly BfYC priorities can be based on branch-specific base priority rankings. This flexibility is one of the advantages of our proposed mechanism.*

Our next Lemma shows that BfYC priorities are compatible with the design of a satisfactory mechanism.

Lemma 2 *BfYC priorities satisfy the unilateral substitutes condition and the law of aggregate demand, and they are fair.*

Lemma 2 implies that COSM is well-defined under BfYC priorities. Let φ^{BfYC} denote the COSM induced by BfYC priorities. This mechanism fixes all previously mentioned shortcomings of the ROTC mechanism:

Proposition 3 *The outcome of φ^{BfYC} is stable under BfYC priorities and it is weakly preferred by any cadet to any stable allocation. Moreover φ^{BfYC} is strategy-proof, fair, and respects improvements.*

Indeed,

Proposition 4 *Given BfYC priorities, φ^{BfYC} is the only mechanism that is stable and strategy-proof.*

We have shown that the potential adoption of the COSM induced by BfYC priorities benefits cadets in numerous ways. Most notably the dead zone is eliminated, restoring the fairness of the mechanism, and the vulnerability of the mechanism to gaming either through preference manipulation or through effort reduction is fully eliminated. In the next section we explain why cadets are not the only beneficiaries of this potential branching reform.

6 Policy Implications for the Army

From a mechanism design perspective, the current ROTC mechanism is a severely deficient mechanism. In this section we show that this is not only a matter of theoretical aesthetics and that the elimination of these shortcomings mitigates several policy problems that the Army has identified.

6.1 Better Utilization of Branch-for-Service Incentives Program

Since 1980s, the U.S. Army has experienced very low retention rates among its most junior officers. This important problem has been well-analyzed, and it is estimated that about 75-80 percent of the required officers at the ranks of Major and Captain are available today (Wardynski, Lyle, and Colarusso 2010). The low retention rates of USMA and ROTC graduates mean that the Army loses much of its ability to screen the quality of its officers for higher ranks. This is evidenced by promotion rates well beyond Defense Officer Personnel Management Act (DOPMA) target rates, as well as the shortened times between promotion opportunities. In contrast to DOPMA target rates of 95 percent and 80 percent, the promotion rates in 2005 to the ranks of Captain and Major were 98.4 percent and 97.7 percent respectively (Henning 2006). Similarly, between 1992 and 2004, the share of captains with less than 4 years of active federal commissioned service rose from 8 percent to 30 percent (Wardynski, Lyle, and Colarusso 2010). To make matters worse, the retention rate of higher quality officers is particularly low, perhaps because they have especially appealing outside opportunities.

The introduction of branch-for-service incentives programs is a response to this problem. The voluntary nature of this program makes it especially appealing. However, restricting cadet bids to only a one-time bid of three additional years reduces the potential impact of the mechanism. Moreover, cadets between 20-50 percent of the OML are to a large extent shut off from the branch-for-service program because of the dead zones they face. They are ineligible for 70 percent of the branch-of-choice slots that are available for bidding. Favoring low-performing cadets at the expense of these moderately well-performing cadets not only undermines the order-of-merit system, but also potentially aggravates their attrition rate. The adoption of the COSM induced by BfYC priorities will not only allow all cadets to bid more than three years for their desired career specialties, it will also allow the Army to distribute talent across branches based on cadet willingness to serve rather than artificially created dead zones. Instead of favoring arbitrary low-performing cadets, our proposed mechanism favors cadets who are most eager to serve in the Army.

6.2 Branch Choice and Diversity among Senior Military Officers

While the enlisted ranks of the U.S. military exhibit a high level of demographic diversity, the leadership of the military has remained demographically homogeneous. In 2006, while 31 percent of the enlisted ranks of the military were African American or Hispanic, only about 16 percent of all officers were African American or Hispanic, and only 5 percent of all Generals were African American or Hispanic (Lim et al. 2009). This is cause for major concern, and significant resources have been devoted to understanding this phenomenon. In a recent Rand Corporation report prepared for the Office of the Secretary of Defense, Lim et al. (2009) conclude that the relative scarcity of minorities in combat arms branches of the Army is a potential barrier to improving demographic diversity in the senior officer ranks. In 2006, 80 percent of all Generals were from combat arm branches. Using 2007 Army ROTC data, Lim et al. (2009) show that while 58 percent of white cadets' submitted first choices were in combat arms, only 31 percent of African American cadets' first choices were in combat arms. They also report that minorities tend to rank lower on the OML and conclude that these numbers may not truly reflect a lack of interest on the part of minorities for combat arms. The following quote is from Lim et al. (2009):

In this exploratory study, we have demonstrated that it is critical for the Army to increase minority representation in key career fields to improve the racial and ethnic diversity of its top military officers. But we also contend that there is a strong need for a more in-depth analysis of the Army branching process. If, as our study suggests, minorities are indeed self-selecting into career fields with relatively limited promotion opportunities, why are they doing so? On the one hand, minority cadets could truly prefer different career fields than white cadets. In this case, policy should focus on ways to make combat career fields more appealing to minorities. On the other hand, minorities may not really prefer support career fields but rather may reason that they lack the OML to get a more competitive career field (or they forecast a low probability of success in that career field). In this case, minority cadets might desire a Combat Arms career field but may opt for their most-preferred

Combat Support or Combat Service Support career field thinking that they would never get a top Combat Arms assignment.

The authors are unable to interpret ROTC preference data because they do not know to what extent minorities strategically avoided more competitive career fields (to avoid a forced assignment). This would not have happened had ROTC used a strategy-proof mechanism. The vulnerability of the ROTC mechanism to preference manipulation thus has adversely affected the authors' ability to prescribe an adequate policy recommendation in this important analysis. There are also several other studies emphasizing the need for understanding minority preferences. The following quote is from Clark (2000):

Another area for future research should focus the issue more closely on the branch selection process in commissioning sources. This issue requires a broad quantitative study to determine the predominant factors in branch selection for black officer candidates from ROTC and USMA. Insight into these factors could lead to inventive solutions to increase ethnic diversity in the combat arms.

These and numerous similar studies show that the adoption of a strategy-proof mechanism is highly valuable to ROTC. Hence even if ROTC is persistent in keeping its current priority structure that relies on dead zones, adoption of COSM will eliminate the difficulties the Army faces in preference data interpretation and allow it to adopt adequate policies to combat minority underrepresentation in its senior ranks.

6.3 Flexibility to Accommodate Branch-Specific Priorities

ROTC leadership currently distributes talent across branches by shutting off the upper-half of the OML from the last 35 percent of slots at each branch. This direct approach heavily relies on the use of a common base priority ranking across all branches. Leadership at some of the branches has been critical of this practice. The following quote is from Besuden II (2008), who argues that the base priorities for Military Intelligence should be improved:

The Army Reserve Officer Training Corps (ROTC) accessions process does not serve either the needs of the Army or the cadets attempting to get one of their top branch

choices. The accessions process overvalues certain aspects of a cadet's background and puts no value on other aspects that could indicate a cadet's potential. In order to get and retain those cadets who are best suited to be Military Intelligence Officers, the accessions process must be changed to better reflect the key competencies of an MI Officer.

As outlined by the United States Army Intelligence Center (USAIC), the MI Corps has three key priorities for its newly commissioned lieutenants. The first is critical thinkers who have the ability to change as their environment changes, the second is technically proficient officers who understand how to apply their craft, and the third is effective communicators that can clearly state their analysis to decision makers. In order for the MI Corps to achieve these goals in their junior officers it needs to bring in officers who are better suited to serve in and succeed in the branch.

Another reason why many are critical about the use of ROTC-OML as the uniform base priority across all branches has to do with the unusually diverse backgrounds of ROTC cadets. Army ROTC is offered in more than 270 universities and colleges across the U.S. Hence the quality of education varies widely across ROTC programs. The weight of academic performance, measured by cadet GPA in the first three years of college, is 40 percent in the OML calculation, and the lack of a standard causes ROTC-OML to be overly subjective.

Our proposed mechanism, unlike the ROTC mechanism, is fully flexible on the choice of base priorities.

6.4 Elimination of the Risk of Cadets Intentionally Lowering OML

The policy implications discussed so far are all on issues heavily debated within the Army. We now present another potential risk of maintaining the ROTC mechanism, even though this risk might not be immediately clear to policy makers in the Army.

Since the ROTC mechanism severely penalizes cadets from the 20th to 50th percentiles of the OML, it gives strong incentives to these cadets to reduce their efforts in their academic, physical fitness, and military studies so that they can be ranked below the median. This incentive is

especially strong for cadets just above the median cadet, since they can avoid losing access to essentially all career branches with a relatively small “compromise” in their OML.¹⁴ While it is hard to know whether cadets actually engage in such forms of manipulation, a mechanism that promotes such behavior can clearly compromise the Army’s efforts in investing its future.

In contrast to the ROTC mechanism, our proposed mechanism respects improvements in cadet performance, and thus cadets can only benefit from an increase in their base priorities. In simple terms, COSM under BfYC priorities fully aligns cadets’ interests with those of the Army.

7 Conclusion

Market design owes much of its recent success to discovery of new practical applications that are supported by elegant theory. Starting in the mid-1990s, auctions have been employed to allocate radio spectrum, electricity, and timber, involving hundreds of billions of dollars worldwide (Milgrom 2004). More recent applications include student admissions (Balinski and T. Sönmez 1999, Abdulkadiroğlu and Sönmez 2003, Ergin and Erdil 2008, Kesten 2010), kidney exchange (Roth, Sönmez, and Ünver 2004, 2005, 2007, and Ünver 2010), course allocation (Sönmez and Ünver 2010 and Budish and Cantillon 2011), and internet ad auctions (Edelman, Ostrovsky, and Schwarz 2007 and Varian 2007). In this paper we have introduced a new practical application of market design. We present a strong case for the replacement of the ROTC mechanism and argue that the Army’s distributional goals can be implemented through a more extensive use

¹⁴The ROTC mechanism is not the only mechanism that harbors incentives for effort reduction. For example, Balinski and Sönmez (1999) show that the mechanism that assigns high school graduates to colleges in Turkey suffers from this deficiency. For most mechanisms, however, expecting agents to materialize such incentives would be unrealistic, for it would require masses of information agents cannot have. This is where the ROTC mechanism stands out, and manipulating it through effort reduction is rather easy. Indeed, the Army provides all the necessary data that is needed for a successful manipulation in the following link: http://www.career-satisfaction.army.mil/pdfs/Order_of_Merit_Score_Calculations.pdf. The data in this document include the order-of-merit score for the median cadet for year 2010 as well as the effect of a 1 point increase in order-of-merit score on the OML ranking around the median.

of market principles. While our focus has been on a potential reform of the ROTC mechanism, our intention is also introducing a resource allocation model where part of the allocation is based on priorities and market principles take over the rest.

Appendix: Proofs

Proof of Lemma 1: Let C_b be the choice function for branch b under ROTC priorities. For $Y \subseteq X$, let $C_b^1(Y), C_b^2(Y), C_b^3(Y)$ denote the set of contracts included in $C_b(Y)$ in Phase 1, Phase 2, and Phase 3 respectively.

1. *ROTC priorities satisfy the unilateral substitutes condition:* Let $X' \subset X$ be a set of contracts and let $x = (i, b, t) \in X'$ be the only contract in X' that involves cadet i . Suppose $x \notin C_b(X')$. Clearly $x \notin C_b^1(X')$, $x \notin C_b^2(X')$, and $x \notin C_b^3(X')$.

Let $X'' \supset X'$. Since $x \notin C_b^1(X')$, there are at least $0.5q_b$ cadets with contracts in X' who have higher OML priority than cadet x_I . Since each of these cadets competes for slots in $C_b^1(X'')$ as well, we have $x \notin C_b^1(X'')$.

Next w.l.o.g assume $x_T = t_2$ for otherwise contract x does not even qualify for slots in $C_b^2(X'')$ or $C_b^3(X'')$. Since the minimum OML priority needed for Phase 1 slots is at least as high under X'' as in X' , any cadet who cannot secure a slot in Phase 1 under X' also fails to receive one under X'' . Hence any contract (with higher cost t_2) that is considered for a Phase 2 slot under X' is also considered under X'' , which implies that the minimum OML priority needed for Phase 2 slots is at least as high under X'' as in X' . Thus $x \notin C_b^2(X')$ implies there are at least $0.15q_b$ higher priority cadets than cadet x_I with higher cost contracts in X'' who fails to receive a slot in Phase 1 under X'' , which in turn implies $x \notin C_b^2(X'')$.

Finally w.l.o.g assume cadet x_I is in the lower half of the OML priority ranking, for otherwise he does not even qualify for slots in $C_b^3(X'')$. Since the minimum priorities needed for Phase 1 and Phase 2 slots are both at least as high under X'' as X' , any contract that is considered for a Phase 3 slot under X' is also considered under X'' . Hence the minimum OML priority needed for Phase 3 slots is at least as high under X''

as in X' and thus $x \notin C_b^3(X')$ implies $x \notin C_b^3(X'')$. Hence $x \notin C_b(X'')$ and therefore ROTC priorities satisfy the unilateral substitutes condition.

2. *ROTC priorities satisfy the law of aggregate demand:* Let $X' \subset X''$. All contracts are eligible for slots chosen in Phase 1 under ROTC priorities. Since every agent who has a contract in X' also has one in X'' , we have $|C_b^1(X'')| \geq |C_b^1(X')|$. Moreover $X' \subset X''$ implies that the minimum OML priority needed for Phase 1 slots is at least as high under X'' as in X' . That means a cadet who cannot secure a slot in Phase 1 under X' also fails to receive one under X'' . Hence any contract (with higher cost t_2) that is considered for a Phase 2 slot under X' is also considered under X'' . Thus $|C_b^2(X'')| \geq |C_b^2(X')|$. It also implies that, as in the case of Phase 1 slots, the minimum OML priority needed for Phase 2 slots is at least as high under X'' as in X' . Finally consider a cadet i who is in the lower half of the OML and suppose $(i, b, t_2) \in C_b^3(X')$. Clearly cadet i does not meet the priority threshold for Phase 1 or Phase 2 slots under X' , and thus under X'' as well. Hence contract (i, b, t_2) is one of the contracts to compete for Phase 3 slots under X'' and hence $|C_b^3(X'')| \geq |C_b^3(X')|$. Since $C_b^1(X''), C_b^2(X'')$ and $C_b^3(X'')$ are disjoint, $|C_b^s(X'')| \geq |C_b^s(X')|$ for $s = 1, 2, 3$ implies $|C_b(X'')| \geq |C_b(X')|$ showing that ROTC priorities satisfy the law of aggregate demand.

◇

Proof of Proposition 1: By Theorem 1, COSM is well-defined when priorities satisfy the unilateral substitutes condition.

Priorities are fair \implies COSM is fair: Fix a choice function C_b for each branch b and suppose that the induced COSM is not fair. Then there exists a problem P and a pair of agents i, j such that

$$\underbrace{\varphi_j(P)}_{=(b,t)} P_i \underbrace{\varphi_i(P)}_{=(b',t')}$$

where $\pi_b(i) < \pi_b(j)$. Observe that cadet i prefers contract $y = (i, b, t)$ to $z = (i, b', t')$. Therefore contract y must be offered to but rejected by the cumulative offer algorithm. Let X' be the set of contracts on hold by the cumulative offer algorithm when y was rejected. Clearly

$y \notin C_b(X' \cup \{y\})$. Furthermore, by choice of X' , cadet i does not have an alternative contract in X' to be considered by the choice function. Let $x = (j, b, t)$ and define $X'' = X' \cup \{x, y\}$. Since priorities satisfy the unilateral substitutes condition, $y \notin C_b(X' \cup \{y\})$ implies $y \notin C_b(X'')$. Moreover, since $\varphi_j(P) = (b, t)$, contract x was picked by choice function C_b when the cumulative offer algorithm has terminated, and therefore again by the unilateral substitutes condition we must have $x \in C_b(X'')$. Hence $x_T = y_T$, $\pi_b(i) < \pi_b(j)$, $x \in C_b(X'')$, and yet $y \notin C_b(X'')$. Finally, by construction, cadet i has no contract other than y in X'' (and thus in $C_b(X'')$ as well). That contradicts the assumption that priorities are fair.

COSM is fair \implies priorities are fair: Suppose priorities are not fair, and for any branch b , let C_b be the resulting choice function. Then there exists a branch b , a set of contracts X' , a pair of contracts $x, y \in X'$ with $x_B = y_B = b$,

$$\left. \begin{array}{l} \pi_b(y_I) < \pi_b(x_I), \\ y_T = x_T, \text{ and} \\ x \in C_b(X') \end{array} \right\} \text{ and yet } y \notin C_b(X').$$

Let $i = y_I$. Construct the following list of preferences P :

1. For cadet i , let pair $(b, y_T) = (b, x_T)$ be the only acceptable pair under P_i .
2. Let $X'' = C_b(X')$. For any agent $j \in X''$ with $x'' = (i, b, x''_T) \in C_b(X')$, let the pair (b, x''_T) be the only acceptable pair.
3. For any remaining cadet, let there be no acceptable branch-cost pair.

Since priorities satisfy the law of aggregate demand, $X'' = C_b(X')$ is the unique stable allocation. Hence $\varphi(P) = X'' = C_b(X')$. But then

$$\underbrace{\varphi_{x_I}(P)}_{=(b, x_T)} P_i \underbrace{\varphi_i(P)}_{=\emptyset}$$

even though $\pi_b(i) < \pi_b(x_I)$. Hence COSM is not fair. ◇

Proof of Lemma 2: Let C_b be the choice function for branch b under BfYC priorities.

1. *BfYC priorities satisfy the unilateral substitutes condition:* Let $X' \subset X$ be a set of contracts and let $x = (i, b, t) \in X'$ be the only contract in X' that involves cadet i . Suppose $x \notin C_b(X')$. Let x^1 be the last contract picked for $C_b(X')$ in Phase 1 and x^2 be the last contract picked for $C_b(X')$ in Phase 2 of construction of $C_b(X')$. We have $\pi_b(x_I^1) < \pi_b(i)$, for otherwise contract x would have been picked for $C_b(X')$ in Phase 1 before contract x^1 . Similarly we have $x_T^2 \geq t$ and if $x_T^2 = t$ then $\pi_b(x_I^2) < \pi_b(i)$, for otherwise contract x would have been picked for $C_b(X')$ in Phase 2 before contract x^2 .

Let $X'' \supset X'$. Let y^1 be the last contract picked for $C_b(X'')$ in Phase 1, and y^2 be the last contract picked for $C_b(X'')$ in Phase 2 of construction of $C_b(X'')$. Since $X'' \supset X'$, the thresholds to be picked are at least as competitive under X'' and hence

- (a) $\pi_b(y_I^1) \leq \pi_b(x_I^1) < \pi_b(i)$,
- (b) $y_T^2 \geq x_T^2 \geq t$ and $y_T^2 = x_T^2 = t \implies \pi_b(y_I^2) \leq \pi_b(x_I^2) < \pi_b(i)$.

Therefore contract x is not chosen for $C_b(X'')$ either in Phase 1 or in Phase 2 showing $x \notin C_b(X'')$. Hence BfYC priorities satisfy the unilateral substitutes condition.

2. *BfYC priorities satisfy the law of aggregate demand:* By construction of the BfYC chosen set, all contracts of a given cadet can be rejected from a branch only when it reaches full capacity. Hence the size of the BfYC chosen set can never shrink as the set of available contracts grows.
3. *BfYC priorities are fair:* Let the set of contracts $X' \subseteq X$ and contracts $x, y \in X'$ with $x_B = y_B = b$, be such that

- (a) $\pi_b(y_I) < \pi_b(x_I)$,
- (b) $y_T = x_T$, and
- (c) $x \in C_b(X')$.

Contract x is picked for $C_b(X')$ either in Phase 1 or in Phase 2. If x is picked for $C_b(X')$ in Phase 1, then $\pi_b(y_I) < \pi_b(x_I)$ implies that the lowest-cost contract of cadet y_I in X'

is picked for $C_b(X')$ in Phase 1 before contract x . Since $y \in X'$, such a contract exists. If x is picked for $C_b(X')$ in Phase 2, $\pi_b(y_I) < \pi_b(x_I)$ and $y_T = x_T$ imply that either the lowest-cost contract of cadet y_I in X' is picked for $C_b(X')$ in Phase 1 or the highest-cost contract of cadet y_I in X' is picked for $C_b(X')$ in Phase 2 before contract x . In either case

$$\exists z \in C_b(X') \text{ such that } z_I = y_I.$$

Hence BfYC priorities are fair.

◇

Proof of Proposition 2: Lemma 1 along with Theorem 1 implies that the outcome of φ^{ROTC} is stable under ROTC priorities and it is weakly preferred by any cadet to any stable allocation. Lemma 1 along with Theorem 2 implies that φ^{ROTC} is strategy-proof. Since cadets in the upper half of the OML priority ranking are denied eligibility for the last 35% of the slots at each branch under ROTC priorities, these priorities are not fair. Therefore φ^{ROTC} is not fair either by Proposition 1. Similarly φ^{ROTC} does not respect improvements since a cadet can gain eligibility for the last 35% of the slots at each branch simply by lowering his OML priority ranking to the lower half.

◇

Proof of Proposition 3: Lemma 2 along with Theorem 1 implies that the outcome of φ^{BfYC} is stable under BfYC priorities and it is weakly preferred by any cadet to any stable allocation. Lemma 2 along with Theorem 2 implies that φ^{BfYC} is strategy-proof.¹⁵ Lemma 2 along with Proposition 1 implies that φ^{BfYC} is fair.

All that remains is to show that φ^{BfYC} respects improvements. Fix a cadet i and let π_1 be an unambiguous improvement for cadet i over π_2 .

Scenario 1: First consider the outcome of φ^{BfYC} under priority order π_1 . Recall that by Remark 1, the order of cadets making offers has no impact on the outcome of the cumulative offer algorithm. Therefore, we can obtain the outcome of $\varphi_{\pi_1}^{BfYC}$ as follows: First, entirely

¹⁵We could alternatively show that BfYC priorities are *substitutable completable*, a condition recently introduced by Hatfield and Kominers (2011), and use Lemma 16, and Theorems 17 and 18 in their paper to prove that φ^{BfYC} is strategy-proof.

ignore cadet i and run the cumulative offers algorithm until it stops. Let X' be the resulting set of contracts. At this point, cadet i makes an offer for his first-choice contract x^1 . His offer may cause a chain of rejections, which may eventually cause contract x^1 to be rejected as well. If that happens, cadet i makes an offer for his second choice x^2 , which may cause another chain of rejections, and so on. Let this process terminate after cadet i makes an offer for his k^{th} choice contract x^k . There may still be a chain of rejections after this offer, but it does not reach cadet i again. Hence cadet i receives his k^{th} choice under $\varphi_{\pi_1}^{BfYC}$.

Scenario 2: Next consider the outcome of $\varphi_{\pi_2}^{BfYC}$, which can be obtained in a similar way: Initially entirely ignore cadet i and run the cumulative offers algorithm until it stops. Since the only difference between the two scenarios is cadet i 's standing in the priority list, X' will again be the resulting set of contracts. Next cadet i makes an offer for his first-choice contract x^1 . Since π_1 is an unambiguous improvement for cadet i over π_2 , precisely the same sequence of rejections will take place until he makes an offer for his k^{th} choice contract x^k . Therefore cadet i cannot receive a better contract than his k^{th} choice under $\varphi_{\pi_2}^{BfYC}$ (although he can receive a worse contract if the rejection chain returns back to him). Hence φ^{BfYC} respects improvements.

◇

Proof of Proposition 4: By Proposition 3, φ^{BfYC} is stable and strategy proof. To show the uniqueness, let φ be a stable mechanism and suppose $\varphi \neq \varphi^{BfYC}$. Suppose φ is strategy-proof. We will show that this assumption results in a contradiction.

Since $\varphi \neq \varphi^{BfYC}$, there exists a preference profile P where $\varphi(P) \neq \varphi^{BfYC}(P)$. Let i be any cadet such that $\varphi(P; i) \neq \varphi^{BfYC}(P; i)$. Since cadet preferences are strict and mechanism φ is stable, we have $\varphi^{BfYC}(P) P_i \varphi(P)$ by Proposition 3. This implies that $\varphi^{BfYC}(P; i) \neq \emptyset$. Let $\varphi^{BfYC}(P; i) = (b, t_\ell)$. Let $P'_i \in \mathcal{P}$ be such that

$$(b, t_1) P'_i \cdots P'_i (b, t_\ell) P'_i \emptyset P'_i (b', t') \quad \text{for any } (b', t') \in (B \times T) \setminus \{(b, t_1), \dots, (b, t_\ell)\}.$$

That is, the only acceptable pairs are (b, t_1) through (b, t_ℓ) under preference relation P'_i .

Since allocation $\varphi^{BfYC}(P)$ is stable under P , it is also stable under (P'_i, P_{-i}) . By Theorem 6 of Hatfield and Kojima, each cadet signs the same number of contracts at every stable allocation

and therefore $\varphi(P'_i, P_{-i}; i) \neq \emptyset$. Hence $\varphi(P'_i, P_{-i}; i) = (b, t)$ where $t \leq t_\ell$. Thus

$$\varphi(P'_i, P_{-i}; i) R_i \underbrace{\varphi^{BfYC}(P; i)}_{=(b, t_\ell)} P_i \varphi(P; i)$$

contradicting the assumption that φ is strategy-proof and completing the proof. \diamond

References

- [1] Abdulkadiroğlu, A., P. A. Pathak and A. E. Roth (2005), The New York City High School Match, *American Economic Review, Papers and Proceedings*, 95, 364-367.
- [2] Abdulkadiroğlu, A., P. A. Pathak, A. E. Roth, and T. Sönmez (2005), The Boston Public School Match, *American Economic Review, Papers and Proceedings*, 95, 368-371.
- [3] Abdulkadiroğlu, A. and T. Sönmez (2003), School Choice: A Mechanism Design Approach, *American Economic Review*, 93, 729-747.
- [4] Balinski, M. and T. Sönmez (1999), A Tale of Two Mechanisms: Student Placement, *Journal of Economic Theory*, 84, 73-94.
- [5] Besuden II, D. L. (2008), Improving ROTC Accessions for Military Intelligence, University of Military Intelligence Report.
- [6] Budish, E. and E. Cantillon (2011), The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard, *American Economic Review*, forthcoming.
- [7] Clark, R.P. (2000), The Lack of Ethnic Diversity in the Infantry: Why Are There So Few Black Infantry Officers in the U.S. Army, Masters Thesis, U.S. Army Command and General Staff College.
- [8] Crawford, V. P. and E. M. Knoer (1981), Job Matching with Heterogeneous Firms and Workers, *Econometrica*, 49, 437-450

- [9] Echenique, F. (2011), Contracts vs. Salaries in Matching, *American Economic Review*, forthcoming.
- [10] Edelman, B., Ostrovsky, M. and Schwarz, M. (2007), Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords, *American Economic Review*, 97, 242-259.
- [11] Ergin, H. and A. Erdil (2008), What's the Matter with Tie-breaking? Improving Efficiency in School Choice, *American Economic Review*, 98(3), June 2008, 669-689.
- [12] Gale, D. and L. Shapley (1962), College Admissions and the Stability of Marriage, *American Mathematical Monthly*, 69, 9-15.
- [13] Hatfield, J.W. and F. Kojima (2008), Matching with Contracts: Comment, *American Economic Review*, 98, 1189- 1194.
- [14] Hatfield, J.W. and F. Kojima (2010), Substitutes and Stability for Matching with Contracts, *Journal of Economic Theory*, 145, 1704-1723.
- [15] Hatfield, J.W. and S.D. Kominers (2011), Contract Design and Stability in Matching Markets, Harvard University and Stanford University working paper.
- [16] Hatfield, J. W. and P. R. Milgrom (2005), Matching with Contracts, *American Economic Review*, 95, 913-935.
- [17] Henning, C.A. (2006), Army Officer Shortages: Background and Issues for Congress, CRS Report for Congress, Congressional Research Service, The Library of Congress.
- [18] Kelso, A. S. and V. P. Crawford (1982), Job Matchings, Coalition Formation, and Gross Substitutes, *Econometrica*, 50, 1483-1504.
- [19] Kesten, O. (2010), School Choice with Consent, *Quarterly Journal of Economics*, 125(3), 1297-1348.

- [20] Kominers, S.D. (2011), On The Correspondence of Contracts to Salaries in (Many-to-Many) Matching, *Games and Economic Behavior*, forthcoming.
- [21] Lim N., J.P. Marquis, K.C. Hall, D. Schulker D., and X. Zhuo (2009), Officer Classification and the Future of Diversity Among Senior Military Leaders A Case Study of the Army ROTC, RAND corporation report.
- [22] Milgrom, M. (2004), *Putting Auction theory to Work*, Churchill Lectures in Economics. Cambridge. Cambridge University Press.
- [23] Roth, A. E. and M. Sotomayor (1990), *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monograph Series, Cambridge University Press, 1990.
- [24] Roth, A. E., T. Sönmez, and M. U. Ünver (2004), Kidney Exchange, *Quarterly Journal of Economics*, 119, 457-488.
- [25] Roth, A. E., T. Sönmez, and M. U. Ünver (2005), Pairwise Kidney Exchange, *Journal of Economic Theory*, 125, 151-188.
- [26] Roth, A. E., T. Sönmez, and M. U. Ünver (2007), Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences, *American Economic Review*, 97-3: 828-851.
- [27] Sönmez, T. and M. U. Ünver (2010), Course Bidding at Business Schools, *International Economic Review*, 51(1): 99-123.
- [28] Sönmez, T. and T. Switzer (2011), Matching with (Branch-of-Choice) Contracts at United States Military Academy, Boston College working paper.
- [29] Sönmez, T. and M. U. Ünver (2010), Matching, Allocation, and Exchange of Discrete Resources, J. Benhabib, A. Bisin, and M. Jackson (eds.) *Handbook of Social Economics*, Elsevier.

- [30] Ünver, U. (2010), Dynamic Kidney Exchange, *Review of Economic Studies*, 77 (1): 372-414.
- [31] Varian, H. (2007), Position Auctions, *International Journal of Industrial Organization*, 25, 1163-1178.
- [32] Wardynski, C., D.S. Lyle, and M.J. Colarusso (2010), *Towards a U.S. Army Officer Corps Strategy for Success: Retaining Talent*, Strategic Studies Institute.

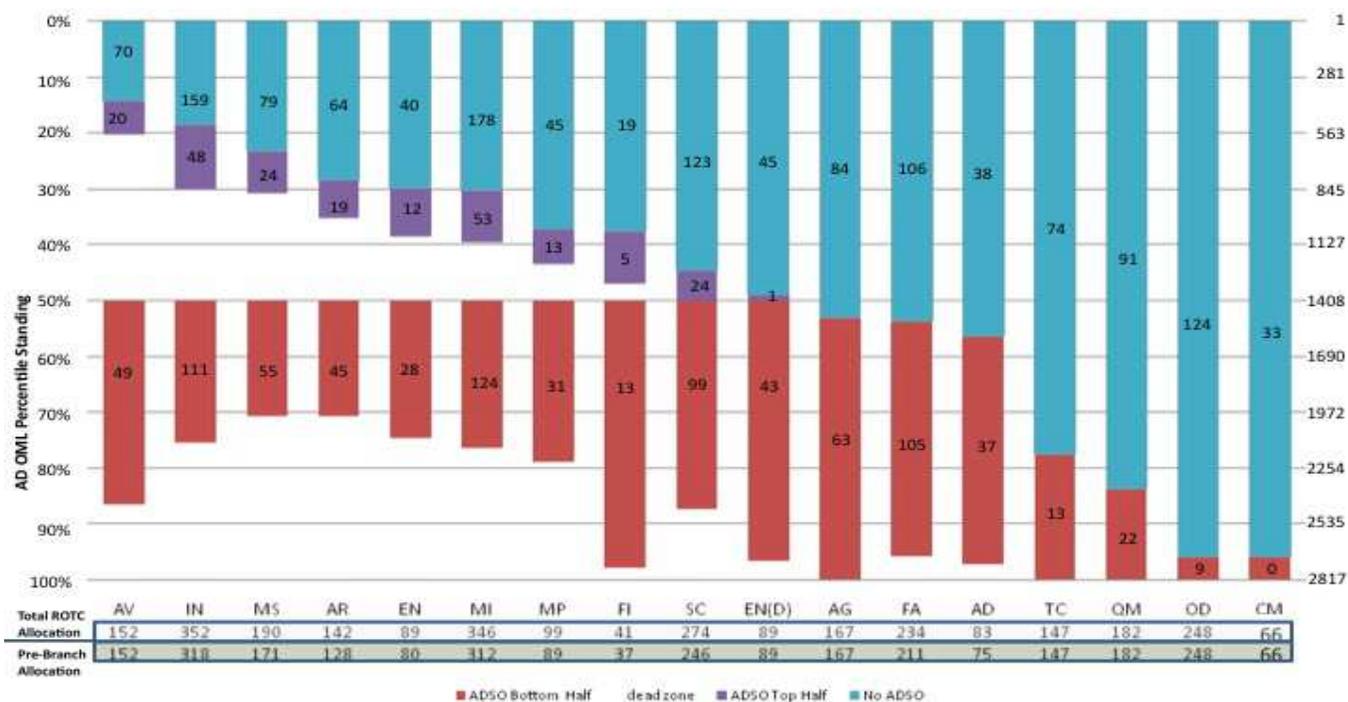


Figure 1: ROTC branch assignment results for year 2011. This graph excludes cadets who are unmatched by the ROTC mechanism and thus received a forced matching by the Department of the Army Branching Board. For each branch, the blue region is that part of the OML standing where cadets secured a slot from the top 50 percent of the slots at the base cost, the purple region is that part of the OML standing where cadets secured a slot between 50-65 percentiles of the slots at the increased cost, and the light brown region is that part of the OML standing where cadets received a slot from the last 35 percent of the slots at increased cost. The blank region between the purple and light brown regions is the dead zone for the most popular eight branches.